Ethical Considerations for Vertically Integrated Hiring Solutions Utilising Machine Learning Models

Richard Day, Yoav Goldberg, Simeon Goldberg

Anthrolytic

March 2026

Abstract

This paper examines the ethical considerations that arise when machine learning models are deployed across the full hiring pipeline, from sourcing to final selection. Rather than proposing a technical or mathematical framework, we focus on the values, obligations, and harms that practitioners must reason about when AI systems mediate access to employment. We argue that responsible deployment requires sustained attention to fairness, transparency, candidate dignity, accountability, privacy, human oversight, and ongoing scrutiny, and that these considerations should shape institutional practice well before any system is built or deployed.

1 Introduction

Hiring decisions are among the most consequential decisions an organisation makes, both for the institution and for the individuals seeking employment. When machine learning systems are introduced into vertically integrated hiring pipelines, the ethical stakes multiply: a single design choice can shape opportunity for thousands of applicants. This paper sets aside questions of model architecture, scoring methodology, and statistical design, and instead asks a prior question: what does it mean to use these systems responsibly? We approach the question from the perspective of practitioners who must defend their choices to candidates, regulators, and the public.

2 Background: AI Hiring and Known Risks

Automated hiring tools have proliferated over the past decade, ranging from resume screeners to video interview analysers. Empirical investigations have repeatedly surfaced concerns: systems that disadvantage applicants on the basis of protected characteristics. These risks are not incidental to the technology; they emerge from the interaction between data drawn from biased social processes and objectives that poorly translate to genuine job performance.

3 Fairness and Non-Discrimination

Fairness in hiring is not a property that can be reduced to a single criterion. Candidates may be wronged through direct discrimination or through evaluative procedures that systematically disadvantage particular groups even when no individual decision is overtly biased. Practitioners deploying AI in hiring should treat fairness as a multi-dimensional commitment, attending to disparate impact across demographic groups, intersectional harms that fall on people belonging to multiple marginalised categories, and the cumulative effects of repeated exposure to similar systems across the labour market.

4 Transparency and Explainability

Candidates have a legitimate interest in understanding how they are being evaluated. Transparency is not merely a regulatory checkbox; it is a precondition for meaningful contestation and trust. Organisations should be willing to disclose that an AI system is being used, what role it plays in the decision, what kinds of signals it considers, and what avenues exist for review. Explainability for affected individuals should be prioritised over post-hoc technical interpretability that is legible only to engineers. Where commercial sensitivity is invoked to limit disclosure, this tension should be acknowledged honestly rather than used as a blanket justification for opacity.

5 Accountability and Governance

Responsibility for hiring decisions must remain locatable. When an AI system influences or determines outcomes, organisations should establish clear lines of accountability that identify who is answerable for the system's design, deployment, and consequences. Vendor relationships should not be permitted to diffuse responsibility into a chain of disclaimers. Internal governance should include named decision-makers, documented review processes, and mechanisms by which affected candidates can escalate concerns to a human with authority to act. Boards and senior leadership should treat hiring AI as a matter of institutional risk and obligation, not solely as an operational efficiency tool.

6 Privacy and Data Protection

Hiring AI systems often process unusually sensitive personal information, including biographical detail, voice and video recordings, behavioural traces, and inferences about character and capability. Ethical practice requires minimising collection to what is genuinely necessary, restricting secondary use, and resisting the temptation to retain candidate data indefinitely on the assumption that it might prove useful later. We assert that, for the vast majority of jobs, the claim that these details are necessary to ascertain their skills is fallicious and often amplifies bias.

7 Auditability and Ongoing Monitoring

Ethical deployment is not a one-time event. Systems drift, populations change, and harms that were invisible at launch may emerge over time. Organisations should commit to ongoing monitoring of outcomes, including disaggregated analysis across demographic groups, periodic external audits, and a willingness to publish findings even when they are unflattering. Audit logs should be maintained in a manner that preserves candidate privacy while enabling legitimate scrutiny by regulators, researchers, and affected parties. The absence of complaints should not be mistaken for the absence of harm.

8 Limitations of This Discussion

This paper does not attempt a comprehensive treatment of every ethical question that hiring AI raises. We have foregrounded considerations we believe are most pressing for practitioners building or deploying such systems, but reasonable people will weight these considerations differently, and additional concerns will surface as the technology and its social context evolve. We have also deliberately avoided proposing a technical or mathematical framework, reserving

such proposals for separate work. The absence of formalism here is intentional: we believe these questions deserve sustained ethical reflection on their own terms before they are operationalised.

9 Conclusion

Vertically integrated hiring solutions built on machine learning models operate at the intersection of opportunity, technology, and power. The ethical obligations they generate cannot be discharged by clever engineering alone. They require institutional commitment to fairness, transparency, candidate dignity, accountability, privacy, meaningful human oversight, and ongoing scrutiny. Practitioners who take these obligations seriously will sometimes find that the most ethical course of action is to deploy less, to deploy more cautiously, or not to deploy at all. We offer this paper not as a final word but as an invitation to deliberate carefully before efficiency arguments are allowed to settle questions that are, at their core, about how we treat one another.